Chapter 5

Synthetic Experiments

In the following Chapters, we present a series of synthetic (or twin) experiments to test our land surface data assimilation algorithm. Synthetic experiments are assimilation runs with synthetically generated parameter, model, and measurement errors, allowing us to compare the estimated and the prior fields to the (synthetic) true fields. Such experiments are ideally suited to evaluate the performance of the algorithm as all of the uncertain inputs are known. Furthermore, synthetic experiments allow us to assess the potential impact of the proposed L-band passive microwave satellite. All experiments are designed to mimic the conditions during the 1997 Southern Great Plains (SGP97) experiment in central Oklahoma (Section 5.3.1) which provides a realistic setup.

In this Chapter, we first explain briefly how synthetic experiments are conducted and explain in more detail why we choose synthetic experiments. Next, we describe the experiment area and the hydrology of the experiment period (Section 5.2). Finally, we describe in detail the data that we use as inputs to the hydrologic model (Section 5.3). The synthetic experiments themselves are described in Chapters 6 and 7.

5.1 Synthetic Experiments and Performance Assessment

5.1.1 Design of Synthetic Experiments

For a synthetic experiment, one realization of the uncertain inputs is generated with a random number generator and suitable mathematical methods to obtain the desired correlation structures. The model is then integrated with this realization of the uncertain inputs, and the output is the corresponding realization of the (synthetic) true state variables. Next, true values of the measured variables are obtained at specified times and locations by applying the measurement equation to the (synthetic) true states. Finally, synthetic measurements are generated by adding synthetically generated measurement error to the true fields of the measured variables.

In the assimilation procedure, the estimation algorithm is denied knowledge of the given (true) realization of uncertain inputs. Instead, it is only supplied with the statistics of the uncertain inputs, that is the prior mean and covariances. The algorithm is also supplied with the noisy synthetic observations, which contain some information about the particular realization of the true fields. From the prior statistics and the noisy synthetic observations, the assimilation algorithm produces a best estimate of the true fields.

If the assimilation algorithm is only supplied with the prior statistics of the uncertain

inputs but not with the synthetic observations, it will default to the so-called prior solution. This prior solution or prior state trajectory is derived from a (forward) model integration with all parameters and model errors set to their prior mean values. In particular, the prior states do not contain any information about the given realization of the true solution. When data are assimilated, the prior solution is the starting point for the algorithm to derive the best estimates of the true states (Section 2.2).

5.1.2 Why Synthetic Experiments?

Synthetic experiments are irreplaceable tools for assessing the performance of the assimilation algorithm. By design, the true solution and the statistics of the uncertain parameters are perfectly known. If we use the same hydrologic model and the same model inputs for the generation of the (synthetic) true fields and for the estimation, the assimilation algorithm is operating under ideal conditions. Non-linear issues aside, the estimate must have certain features that can be tested with the hypothesis tests described in Sections 2.3.6 and 2.4.1. Consequently, coding errors can easily be detected and the effects of nonlinearities can be assessed. Moreover, observing system characteristics can be evaluated and optimized. In Chapter 6, we describe a series of assimilation experiments under such ideal conditions.

When field data are assimilated, the hydrologic model will only be a crude approximation of the "model" that nature is using, and the error statistics we specify will likely be poor approximations of the true error characteristics. By specifying different parameters or statistics in the estimation process than have been used for the generation of the synthetic uncertain inputs, such nonideal situations can be investigated. With synthetic experiments, we can therefore investigate the sensitivity of the estimation algorithm to wrongly specified error statistics.

Similarly, inputs such as the soil hydraulic parameters will only be poorly known. By using different soil hydraulic parameters for the generation of the (synthetic) true fields and for the estimation process, we can emulate the realistic condition of assimilating data into a model that contains structural errors and bias. To compensate for the discrepancy, we will likely have to increase the model error variance in the assimilation algorithm. If we conduct synthetic experiments under such nonideal conditions, we can investigate the influence of the quality of the hydrologic model on the assimilation. Synthetic experiments allow for the evaluation of any number of such scenarios.

A final word of caution is in order. We believe that synthetic experiments are useful and irreplaceable tools for testing the assimilation algorithm and for investigating the sensitivity of the estimation process to various factors. Synthetic experiments cannot, however, replace the ultimate test of the algorithm with field observations. Such a field application must eventually be carried out.

5.2 Experiment Area and Period

The area of the synthetic experiments is located within the SGP97 experiment area. We choose the domain to coincide loosely with the swath of the airborne ESTAR brightness temperature measurements (Section 5.3.1). Figure 5.1 shows the horizontal grid of 16 by 32 estimation pixels together with the county lines and the meteorologic stations of the Oklahoma Mesonet. At a resolution of $5km \times 5km$, the experiment area covers a total of $12,800km^2$ in an $80km \times 160km$ rectangle. The estimation pixels are numbered from 1 to

512 starting in the southwestern corner and going north until the boundary of the domain is reached. The numbering continues at the bottom of the next column to the east of the first column and so on until the northeastern corner of the domain is reached.

The vertical resolution of the saturation is as shown in Figure 4.1 (Section 4.1). Each column has $N_z = 7$ vertical nodes for the saturation W_g at $0cm$, $-5cm$, $-15cm$, $-30cm$, $-45cm$, $-60cm$, and $-90cm$. There is only one node per column for each of the other five states (canopy temperature T_c, vapor pressure e_a, temperature in the canopy air space T_a, soil temperature T_g, and interception water content W_c). Per pixel, we have therefore 12 states, and the state vector has a dimension of $512 \cdot 12 = 6144$ at every time step. (Note that the total number of scalar data in the reference experiments is also 6144 by coincidence.) The total number of 15 minute time steps in the synthetic experiment is 1280 (see below), and the total number of states is therefore $6144 \cdot 1280 \approx 10^7$.

All experiments extend over a two-week period from June 18, 1997 (day 169) to July 2, 1997 (day 183). Figure 5.2 shows time series of the area average micro-meteorologic inputs as derived from the Oklahoma Mesonet data (Section 5.3.2). The top panel shows the area average precipitation. Initially, there is a four day drydown across the entire domain. On day of year 174, significant rain is falling across the entire area. In contrast, the two major precipitation events of days 177 and 179 are concentrated in the northern half of the domain. The total rain over the two-week period measured at the stations in the northern half was between $3cm$ and $8cm$, whereas only $0.5cm$ to $2cm$ were observed at the southern stations. The area average cumulative precipitation over the two-week period is $2.8cm$.

The other panels of Figure 5.2 show the area average air temperature, wind speed, incoming shortwave radiation, and relative humidity. These fields are fairly homogeneous across the experiment area. Note that the area average quantities are only shown for illustration. In the synthetic experiments we use spatially distributed data. For details see Section 5.3.2.

5.3 Inputs to the Hydrologic Model

In this Section, we briefly describe the sources and the character of the various inputs to the hydrologic model that we use in the synthetic experiments.

5.3.1 The Southern Great Plains 1997 (SGP97) Hydrology Experiment

The Southern Great Plains 1997 (SGP97) Hydrology Experiment took place in the sub-humid environment in Oklahoma over the one-month period of June 18 – July 17, 1997 [Jackson, 1997; Jackson et al., 1999]. Its main objectives are

1. to examine the estimation of surface soil moisture and temperature using remote sensing at a hierarchy of scales,

2. to examine the feasibility of estimating vertical profiles of soil moisture and temperature by combining in situ data, remote sensing measurements at the surface, and modeling techniques,

3. and to evaluate the influence of soil moisture on the local surface energy budget and the influence of mesoscale variability in the surface energy budget on the development of the convective boundary layer.

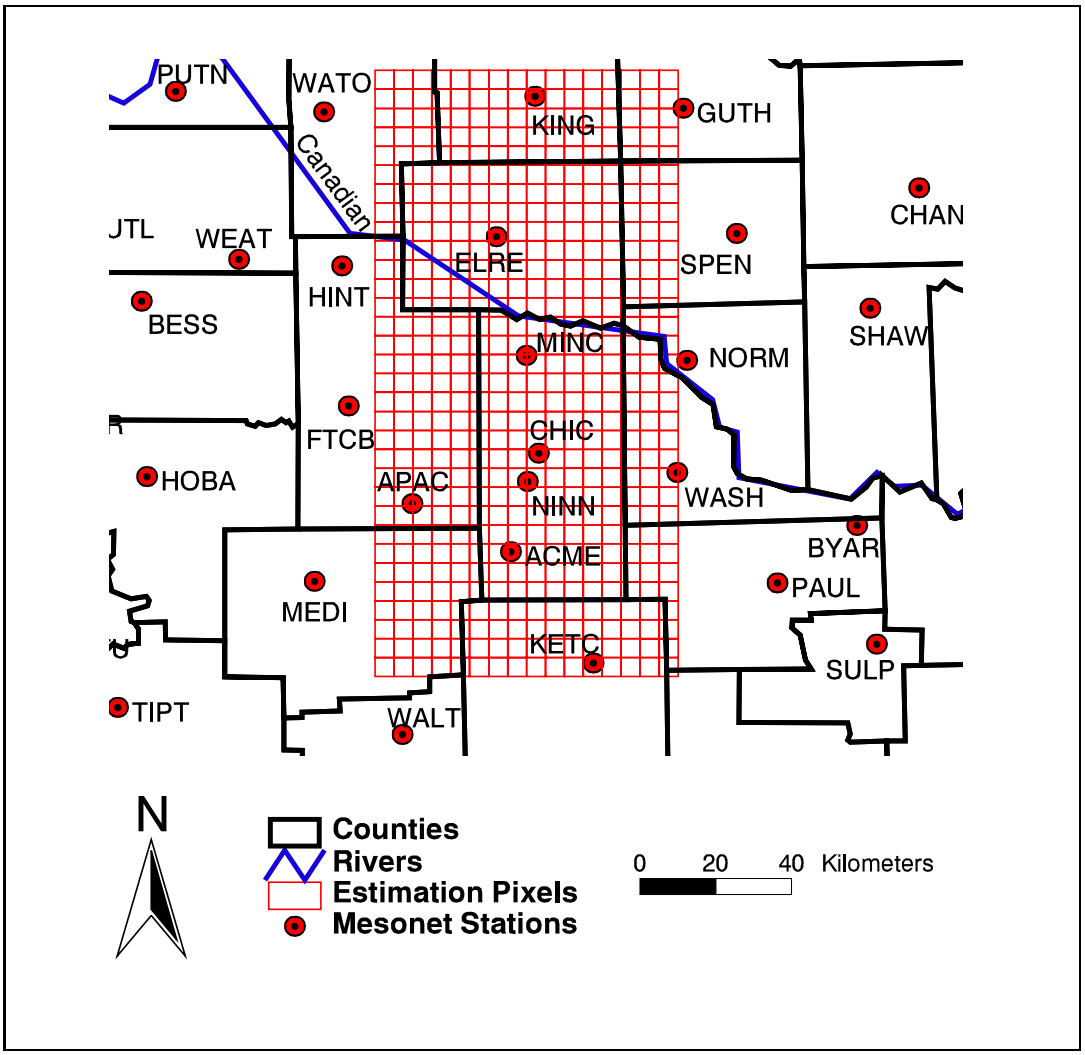

Figure 5.1: Area for the synthetic experiment. The horizontal grid of 16 by 32 pixels ($80km \times 160km$) is shown together with the locations of the Oklahoma Mesonet stations. The horizontal resolution is $5km \times 5km$. Oklahoma City is just to the north of Mesonet station NORM, on the eastern boundary of the domain.

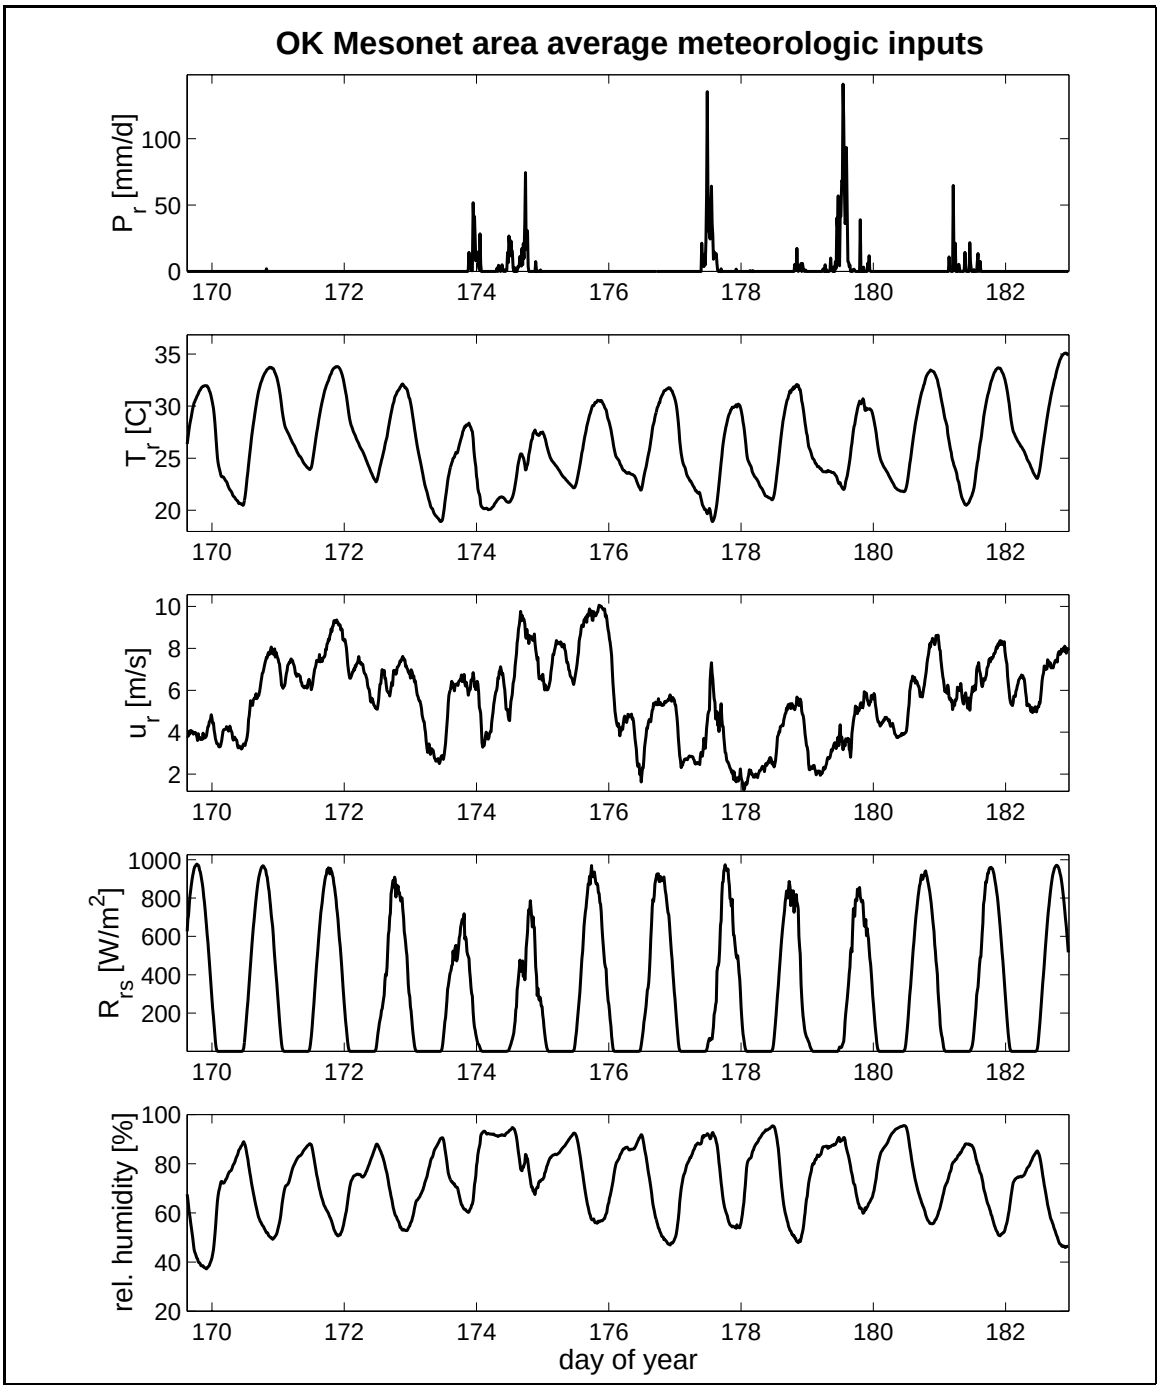

Figure 5.2: Area average micro-meteorologic inputs as derived from the Oklahoma Mesonet data. The panels show (from top to bottom) the area average precipitation, air temperature, wind speed, incoming shortwave radiation, and relative humidity. The two major precipitation events are concentrated in the northern half of the domain. The area average quantities are shown for illustration only. In the experiments, we use spatially distributed data.

To achieve these objectives, a host of ground-based and remotely sensed data were gathered. This includes measurements of ground-based soil moisture and soil properties, airborne boundary layer measurements, atmospheric soundings, vegetation parameters, airborne remote sensing measurements, ground-based remote sensing observations, satellite remote sensing measurements, surface flux data, and surface hydrometeorologic measurements. Moreover, micro-meteorologic data were provided by the Oklahoma Mesonet (Section 5.3.2), and soil texture classes were compiled by the Earth System Science Center (ESSC) at Pennsylvania State University (Section 5.3.3).

In the synthetic experiments of Chapters 6 and 7, we use the Oklahoma Mesonet data for the micro-meteorologic inputs, the ESSC database for the soil texture classification, and the land cover map to obtain the vegetation class for each pixel. In addition, for the Radiative Transfer model we use the maps of the vegetation water content W_v, the surface roughness parameter β^{ε_g}, and the vegetation parameter β^{δ_c}.

The land cover data are derived from Landsat Thematic Mapper (TM) images and come in $30m$ resolution. The parameters of the Radiative Transfer model are available at a resolution of $800m$. According to their type, the data were aggregated to the resolution of the estimation grid by assigning the class of the block majority to the estimation pixel or by averaging over the estimation pixels. Figure 5.3 shows a map of the land cover classes aggregated to the $5km$ resolution. The projection is Universal Transverse Mercator (UTM), Zone 14, which is used throughout this thesis. The predominant land cover class is pasture. In the northern half, there are also large areas of wheat cultivation. The remaining land cover classes include forage, shrub, urban, and water. For the synthetic experiment, we neglect the latter two classes and use the pasture properties for the corresponding pixels as well as for the pixels for which there are no data available. This change does not affect the results of the synthetic experiments. When field data are assimilated, however, the water bodies and urban areas need to be treated separately.

Remotely Sensed Brightness Temperature from ESTAR

During the SGP97 field campaign, the Electronically Scanned and Thinned Array Radiometer (ESTAR) was flown daily on the NASA P3 aircraft. ESTAR successfully recorded sixteen images of L-band brightness temperatures ($1.4GHz$) on a swath of roughly $50km \times 200km$ at a resolution of $800m$. Together with the ground-based soil moisture [Famiglietti et al., 1999] and flux measurements, these invaluable data will allow a first test of the soil moisture assimilation algorithm on field data.

5.3.2 Micro-meteorologic Inputs: The Oklahoma Mesonet

The micro-meteorologic inputs to our land-surface model consist of six data types: precipitation P_r, incoming shortwave radiation R_{rs}, air temperature T_r, vapor pressure e_r, wind speed u_r, and depth average soil temperature T_d. All of the above are directly available or easily derived from the Oklahoma Mesonet data [Brock et al., 1995].

The Oklahoma Mesonet is a statewide network of 115 automated weather observing stations. It is unique in its dense spatial coverage, with an average distance between sites of about $31km$ [Basara et al., 1999]. Among other parameters, each site records at five minute intervals relative humidity at $1.5m$ height above the ground, air temperature at $1.5m$, average wind speed at $10m$, rainfall, barometric pressure, and solar radiation. In

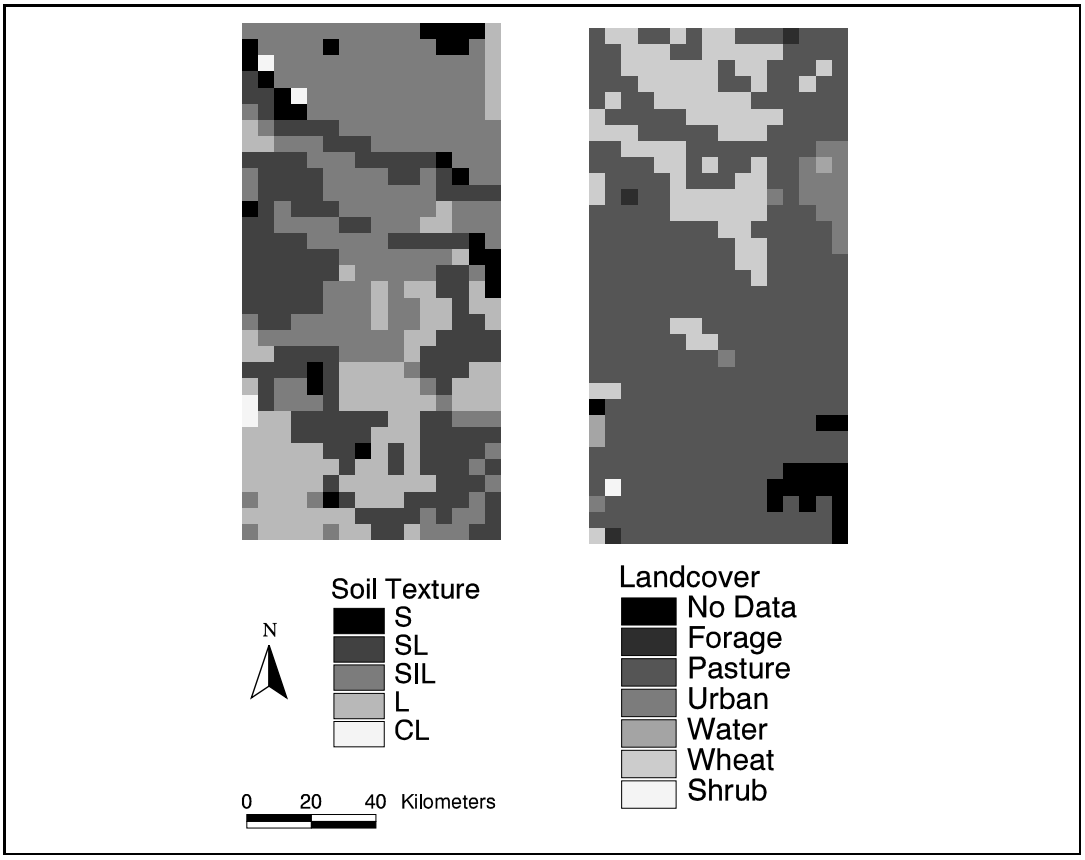

Figure 5.3: Surface soil texture and land cover classes for the synthetic experiment. The surface soil texture map has been derived from the ESSC data base (Section 5.3.3), and the land cover map has been obtained from SGP97 data (Section 5.3.1). The dominant soil texture classes are sand loam (SL), silt loam (SIL), and loam (L). There are also some areas where sand (S) and clay loam (CL) can be found. The dominant land cover class is pasture. In the northern half of the domain, there are also large areas of wheat. The urban, water, and no-data pixels are treated as pasture in the synthetic experiment. The horizontal resolution is $5km \times 5km$.

addition, bare and vegetated soil temperature at $10cm$ depth are measured at 15 minute intervals.

There are a number of supplemental parameters available. Of most interest to us are observations of the bare and vegetated soil temperatures at $5cm$ and at $30cm$, which are taken at nearly half of the stations. Recently, many stations have been upgraded with heat dissipation sensors and time-domain reflectometry (TDR) probes to record soil moisture profiles [Basara et al., 1999]. Unfortunately, at the time of this writing there were still a few unresolved issues in the calibration of the moisture probes. In the future, such operational ground-based soil moisture data have great potential for assimilation into a land-surface model and for the validation of such algorithms.

Preprocessing the Micro-meteorologic Inputs

Before using the Mesonet data as inputs to our land-surface model, we apply a few preprocessing steps. The very few missing data points are filled in with neighboring observations or earlier observations at the same site. The observed relative humidity is converted into the vapor pressure using the measured air temperature. The deep soil temperature of the force-restore equation (3.7) is obtained from a monthly average of the air temperature.

The five minute interval for the micro-meteorologic inputs is attractive from a modeling point of view, as it allows to resolve the dynamics very finely. On the other hand, the model runs become very computationally demanding. We therefore average the five minute data to 15 minute inputs. Consequently, the basic time step of the tangent-linear and the adjoint models is 15 minutes. Note, however, that the initial prior state trajectory is obtained by solving the nonlinear land surface model with a variable time step (Section 4.10). During and after a rain event, the time step of the nonlinear model is typically reduced to a few seconds.

Even though the network of the Oklahoma Mesonet is very dense, the resolution is still too coarse for direct use in the hydrologic model. We interpolate the data with inverse square distance weights to the grid of estimation pixels. The weights for the seventeen stations of the synthetic experiment are shown in Figure 5.4.

It is important to note that we did not place emphasis on the optimal preprocessing of the micro-meteorologic data. Our focus is on the assimilation algorithm rather than on the calibration of the hydrologic model. Most importantly, the quality of the meteorologic inputs will be a lot poorer in future operational applications of land-surface hydrologic data assimilation. By including model error, the data assimilation algorithm is expressly designed to account for such deficiencies (Section 4.5).

5.3.3 Soil Properties

The database of the Earth System Science Center (ESSC) at Pennsylvania State University contains a variety of geographically referenced data sets. For the convenience of SGP97 researchers, the ESSC derived soil properties and land cover data from the State Soil Geographic (STATSGO) data set compiled by the Natural Resources Conservation Service (NRCS) of the U.S. Department of Agriculture. In our synthetic experiments, we use the dominant soil texture data, the sand and clay fraction data, and the bulk density data, all of which are available at $1km$ resolution for each of 11 standard soil layers.

The data sets come in the Universal Transverse Mercator (UTM), Zone 14 projection, which we use throughout this study. The aggregation of the data to the estimation grid

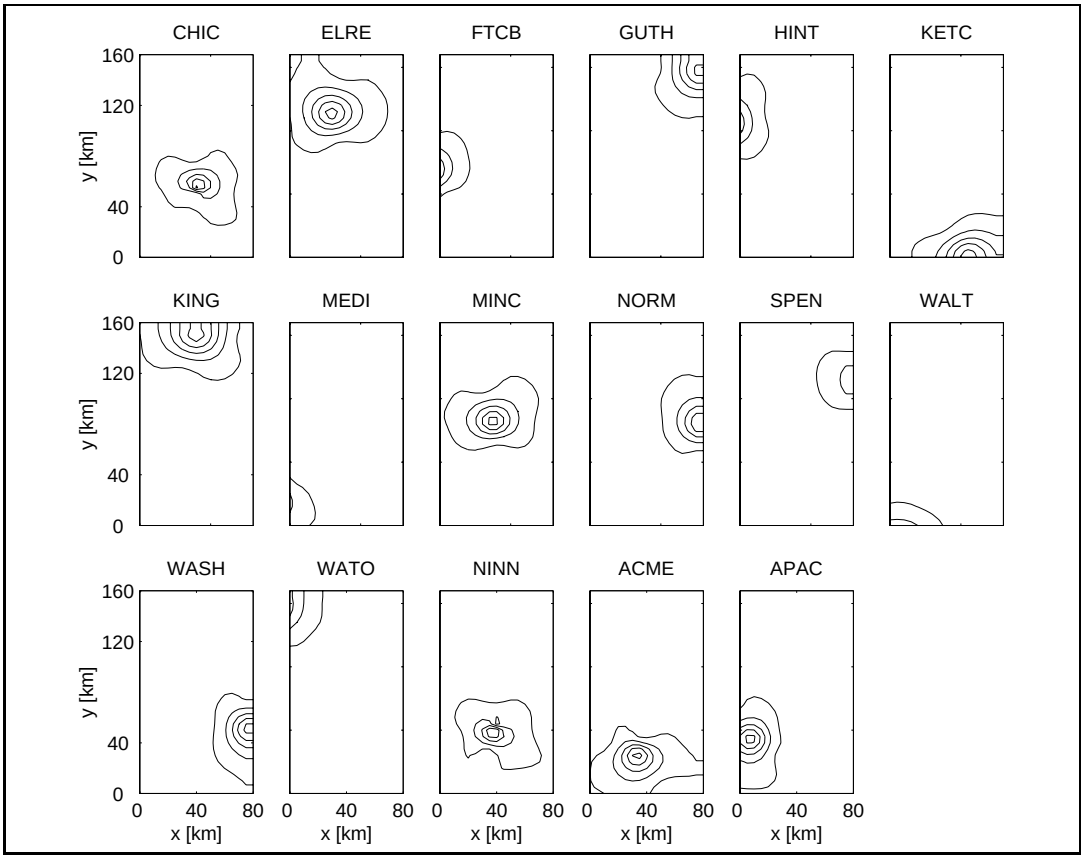

Figure 5.4: Inverse square distance weights for the interpolation of the micrometeorologic inputs. The 0.1, 0.3, 0.5, 0.7, and 0.9 contour lines are shown. The locations of the seventeen Oklahoma Mesonet stations used in the synthetic experiments of Chapters 6 and 7 are also shown in Figure 5.1. The average distance between stations is only $31km$.

was performed according to the data type by assigning the class of the block majority or by averaging. Figure 5.3 shows a map of the surface soil texture classes at $5km$ resolution. The dominant textures are sand loam (SL), silt loam (SIL), and loam (L). There are also some spots with sand (S) and clay loam (CL). Note the streaky pattern of both the soil texture and the land cover classes in the northwestern corner of the domain. This pattern can be recognized in the soil moisture fields during drydowns (Figure 6.3).

5.3.4 Other Data Sources

Not all of the inputs necessary to run the hydrologic model are provided by the SGP97 data set. Each soil texture and land cover class still needs to be translated into the corresponding numerical parameters. Vegetation parameters for the Little Washita catchment, which is located within the SGP97 experiment area, are tabulated in [Kustas and Jackson, 1999]. Inputs to land-surface schemes within General Circulation Models can be found in [Sellers, Los, Tucker, Justice, Dazlich, Collatz and Randall, 1996]. Similar data for the roughness length z_0, the fractional vegetation cover f_c, the vegetation albedo a_c, the minimum stomatal resistance $r_s^{\min}$, and the leaf and stem area indices LAI and SAI can be found in [Dickinson et al., 1993]. Typical values for the roughness length z_0 are also tabulated in [Dorman and Sellers, 1989]. Yang et al. [1998] give typical rooting depths. The soil hydraulic parameters are based on the data published by Clapp and Hornberger [1978]. Finally, the parameters β^{r_c} and β^{r_d} of the resistance network are given in [Sellers and Dorman, 1987].

The choice of the above parameters amounts to calibrating the hydrologic model, a task which we performed only in a rudimentary fashion. Our focus is on the performance of the assimilation algorithm and not on model calibration. For the synthetic experiments presented here the calibration matters little. However, for the future assimilation of field data such as the ESTAR brightness observations, the model must be reasonably well calibrated.